

05-30-00

A

Attorney Docket: 59322

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Patti
For: Dynamic Configuration of Storage
Arrays

CERTIFICATE OF MAILING BY "EXPRESS MAIL"

"Express Mail" Mailing Label Number

EK334213039US

Date of Deposit

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" Service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231

Calvin B. Ward

Signature of person mailing

1c525 U.S. PTO
09/580936
05/25/00

PATENT APPLICATION

Hon. Commissioner of Patents and Trademarks
Washington, D.C. 20231

Sir:

The following documents are transmitted herewith:

(X) original patent application. (X) Utility () Design () continuation-in-part application.

(X) The Declaration and Power of Attorney. (X) signed () unsigned or partially signed
(X) 3 sheets of (X) formal drawings () informal drawings
() Small Entity Statement

(X) A check in the amount of \$690 to cover the filing fee is enclosed.

Applicants believe that no additional fee is required. However, the Commissioner is hereby authorized to charge any additional fees which may be required in this application under 37 C.F.R. Section 1.16-1.17 during its entire pendency, or credit any overpayment, to Deposit Account No. 23-0424. Should no proper payment be enclosed herewith, as by a check being in the wrong amount, unsigned, post-dated, otherwise improper or informal or even entirely missing, the Commissioner is authorized to charge the unpaid amount to Deposit Account No. 23-0424. This sheet is filed in duplicate.

Calvin B. Ward
Registration No. 30,896

18 Crow Canyon Court, Suite 305
San Ramon, CA 94583
Telephone: (925) 855-0413
Telefax: (925) 855-9214

Dynamic Configuration of Storage Arrays

Field of the Invention

5 The present invention relates to memory systems, and more particularly, to a memory system which detects errors and reconfigures itself to avoid bad memory cells.

Background of the Invention

10 As the cost of computational hardware has decreased, computers with ever-larger memory systems have proliferated. Systems with hundreds of Mbytes are common, and systems with a few Gbytes of memory are commercially available. As the size of the memory increases, problems arising from bad memory cells become more common.

15 Memory failures may be divided into two categories, those resulting from bad memory cells that are detected at the time of manufacture and those that arise from cells that fail during the operation of the memory. At present, problems arising from defective memory cells that are detected during the manufacturing process are cured by replacing the bad cells. The typical memory array is divided into blocks. Each memory chip has a predetermined number of spare
20 blocks fabricated thereon. If a block in the memory is found to have a defective memory cell, the block in question is disconnected from the appropriate bus and one of the spares is connected to the bus in its place. However, once the part is packaged, there is no means for replacing a block with a spare, since the replacement process requires hard wiring of the spares to the bus.

25 The cost of testing the memory chips is a significant factor in the cost of the chips. The rate at which memory cells can be tested is limited by the internal organization of the memory blocks and the speed of the buses that connect the memory blocks to the test equipment. The various buses are limited to speeds of a few hundred MHz. Data is typically written and read as blocks having 64 bits or less. Since a write operation followed by a read operation requires
30 several clock cycles, the rate at which memory can be tested is limited to 100 million tests per

second. Extensive testing requires each memory cell to be tested a large number of times under different conditions such as temperature and clock speed. Hence, a 1 Gbyte memory chip would require minutes, if not hours, to thoroughly test. The cost of such testing would be prohibitive; hence, prior art memory chip designs will not permit extensive testing at the 1 Gbyte level and beyond.

Even when the obviously bad memory blocks have been removed, sooner or later, the memory will fail because of the failure of one or more cells in a block. The probability that such a failure will cause a system failure depends on the lifetime of the system, the size of the memory, and the type of memory. The probability of such a failure increases with the lifetime of the system and the size of the memory. While system lifetimes are not increasing, the size of memory is increasing. Accordingly, more system failures are expected.

In addition, some types of memory cells have higher failure rates than others. For example, EEPROM and flash memories can only be written a relatively small number of times compared to conventional DRAM and static RAM memories. In the case of EEPROMs and flash memories, the limited number of write cycles imposes severe restrictions on the possible applications of these memories. Similarly, memories based on ferroelectrics have relatively small lifetimes relative to these conventional memories; however, the ferroelectric memories can be written many more times than EEPROMs and flash memories.

In principle, all of these types of memories would benefit by having some form of reconfiguration system built directly into the memory. Such a system would replace blocks of memory that fail during the operational life of the system, thereby extending the lifetime of the system. However, prior to a memory cell actually failing, there is often a period of time in which the memory cell operates, but with a high error rate. Such a memory cell can cause intermittent system failures and may be very difficult to diagnose. Hence, any form of block replacement system that depends on detecting the failure of a block may not be able to operate successfully.

Broadly, it is the object of the present invention to provide an improved memory system.

coupled to that bit line when that row address is coupled to the row select circuit. The memory also includes a word assembly circuit for selecting N bit lines from said M bit lines, where N is less than or equal to M. The word assembly circuit includes a memory for storing a mapping specifying the N bit lines for each possible row address. In such embodiments, the controller
 5 alters the mapping to eliminate a reference in the mapping to a bit line that causes a defective storage cell to couple data to a bit line in response to one of the row addresses. The memory also includes an error correcting circuit for detecting errors in data words. The error correcting circuit generates a corrected data word and an error data word from the N data values coupled thereto, the error data word indicating which of the N data values, if any, was erroneous. The N data
 10 values are determined by a word assembly circuit that connects N of the M bit lines to the error correcting circuit, where N is less than or equal to M. A control circuit connected to the error correcting circuit uses the error data words and the row addresses to alter the mapping in the row select circuit in response to the error data words so as to avoid defective memory storage cells or bit lines.

Brief Description of the Drawings

Figure 1 is a block diagram of a memory 10 according to the present invention.

Figure 2 is a block diagram of a memory 50 that utilizes an error correcting codes
 20 according to the present invention.

Figure 3 is a block diagram of a memory 100 constructed from a plurality of memory banks according to the present invention.

Detailed Description of the Invention

The manner in which the present invention gains its advantages can be more easily understood with reference to Figure 1 which is a block diagram of a memory 10 according to the present invention. For the purposes of this discussion, it will be assumed that memory 10 stores
 30 data words having N_D bits per word. The data is stored in a two-dimensional array 11 of memory

cells that are organized into a plurality of rows and columns. A typical storage cell is shown at 15. All bits of any given data word are stored on the same row. The specific row is determined by a mapping that is stored in a content addressable memory (CAM) 43 that stores a mapping of the logical addresses used by the outside computer system to the row addresses used with memory 10.

In the preferred embodiment of the present invention, a number of words are stored in each row. A word is selected via a row select circuit 14 which causes the storage cells associated with that row to be connected to a plurality of bit lines. Hence, by placing the appropriate signal on line 13, storage cell 15 is connected to bit line 12. If data is to be written into the word, the data values are coupled to the appropriate bit lines via a cross-connect switch 20. If data is to be read from the word, the data values stored in the storage cells on the selected row are placed on the bit lines by the selected storage cells and decoded by a block of sense amplifiers 21. The specific bit lines corresponding to the desired word are then selected by cross-connect switch 23. The mapping of the physical bit lines to the logical address used by the outside computer system is stored in a row CAM 43.

There are additional rows and columns of storage cells in array 11 beyond those needed to store the number of bits for which the memory was designed. These spares can be used to replace the defective memory cells by changing the mapping stored in CAMs 42 and 43. In one embodiment of the present invention, each time memory 10 is powered up, a controller 40 loads the CAMs with default values and tests the memory array to determine if any bad rows or columns exist. If a bad row is found, controller 40 alters the mapping in CAM 43 such that the bad row is replaced by one of the spare rows. Similarly, if a bad column is found, the mapping in CAM 42 is altered such that the bad column is replaced by one of the spare columns.

As noted above, test speed has been a problem in prior art memory designs. The present invention overcomes the speed limitations by providing the testing function on the memory chip itself and by using a highly parallel testing strategy. During the testing phase, controller 40 stores and retrieves test values from entire rows. In general, the sense amplifier block 21

includes a register having one bit per bit line that latches the value read from the currently selected row. The values in this register may be connected back to the bit lines during the testing operation. Hence, many words can be tested in a single write/read operation. For example, in a 1 Gbyte embodiment of the present invention, each row includes 72 Kbits of storage.

Furthermore, the speed with which the controller can perform a test cycle on a row is not limited by the speed of the external bus that connects the memory to the central processing unit (CPU) in the computer system. Accordingly, the present invention can run at higher internal clock speeds during testing, and hence, provide even faster memory testing.

Finally, memory 10 is typically one of several memory blocks on a memory chip. For example, the 1 Gbyte embodiment of the present invention described above includes 4 such blocks. During testing, each block is isolated from the others and can perform tests in parallel with the other blocks. The combination of the high degree of parallel processing and increased internal speed allow the present invention to test memory cells at the rate sufficient to identify even intermittently failing cells prior to bringing the memory on line. The memory can then be configured to prevent data storage in the bad cells.

In the preferred embodiment of the present invention, the actual data stored in the storage array is encoded via an error correction code to further improve the reliability of the memory. Refer now to Figure 2, which is a block diagram of a memory 50 according to the present invention that utilizes an error correcting code. To simplify the following discussion, those elements of memory 50 that serve the same functions as elements shown in Figure 1 have been given the same numerical designations and are assumed to perform the same functions unless the text indicates otherwise.

In memory 50, each word also includes a number of error correcting bits that are stored with the word. The error correcting bits are added to the data word by an ECC generator 30 prior to storing the combined data and error correcting bits in storage array 11. Hence, each word actually occupies N storage cells, where $N > N_D$.

The error correcting bits serve two functions. First, the error correcting bits allow some bit errors to be corrected by error correction circuit 31. In the preferred embodiment of the present invention, sufficient error correcting bits are provided to allow all one-bit errors to be corrected and all two-bit errors to be detected. Error correcting codes that provide this type of functionality are well known to those in the art, and hence, will not be discussed in detail

Second, the error correcting bits allow the present invention to determine which bits failed. This information is sent to controller 40 by error correcting circuit 31. Controller 40 keeps track of the locations of the errors to determine if one or more rows or columns of storage cells, or the electronics associated therewith are failing. When controller 40 determines that a particular column or row is failing, controller 40 substitutes a spare column or row included in storage array 11 for the failing column or row by altering the appropriate mapping in either CAM 42 or 43. In effect, memory 50 is in a test mode even when it is operating to store data.

Column replacements are much more costly, both in terms of down time for the memory and hardware, than row replacements. A row replacement can be accomplished in a single read-write cycle, since all of the bits of the old row are available at the same time and all of the bits of the new row can be written at once using the register in the sense amplifier block discussed above. In memories based on conventional DRAMs, the row replacement can be carried out as part of the normal refresh cycle.

The replacement of a column is somewhat more complicated, since each row in the defective column must be read and rewritten such that each cell in the defective row is replaced by a corresponding cell in one of the spare columns. During such a re-configuration involving a row replacement, the memory may not be available for a significant period of time. During the time a column is being replaced, the present invention stalls the central processor by making use of the serial master (SM) bus that is incorporated in many processor designs for out of band communications between memory and the CPU.

In addition, column replacements require the two cross-connect switches described above; while row replacements require only the register in the sense amplifier block. To reduce the complexity of these switches, the present invention utilizes a more restrictive form of column replacement. In the preferred embodiment of the present invention, each row can be replaced by only a limited number of the available spare rows. For example, a row might be defined to have 2 spare columns of memory cells between each two bytes on the row. These spares would service the bytes immediately adjacent to the spares; hence, the cross-connect switch would need only 32 switches for connecting these two bytes to the two spares. In contrast, 288 switching elements would be needed if all possible columns are to be used in any bit position. Hence, this restriction reduces the number of switching elements in the cross-connect switches, since a switching element is only needed between the lines that could potentially be connected to one another.

If the number of data bits per word is large compared to the number of bits in the address, it is more efficient to use a single bit CAM 60 for the spares than to replace an entire row when one bit in the row is defective. In embodiments of the present invention that utilize such single bit replacements, the data and bit location are stored in CAM 60 which inserts the data bit into the outbound data word with the aid of a multiplexer 61. Data for the defective bit is copied into CAM 60 at the same time the data for the remaining bits is stored in storage array 11. The specific bits that are stored in CAM 60 are determined by controller 40. To simplify the drawing, the connections between controller 40 and CAM 60 have been omitted.

In general, there will be a plurality of storage arrays such as storage array 11 that share the same cross-connect switches, sense amplifiers, etc. For each such array, column CAM will provide one mapping of the columns to the bits of the data words. The identity of the block currently connected to the sense amplifiers, etc. is determined by the row address. The blocks may be viewed as a partitioning of the rows into a sequence of blocks that share the same sense amplifier block; hence, the loss of a sense amplifier may still require a column replacement. Each block, however, has its own bit lines that operate independently of the bit lines of the other

blocks; hence, a break in a bit line can be accommodated by mapping the rows to other blocks rather than replacing an entire column.

As noted above, column replacements are quite costly compared to row replacements; hence, it should also be noted that systems that do not use column replacements may be constructed without deviating from the teachings of the present invention. In such systems, the column CAM and much of the cross-connect hardware can be eliminated. The column CAM allows one to replace one column of storage cells with another column of storage cells. It provides the greatest benefit when one sense amplifier fails. If there are sufficient error correcting bits to correct for several bits in each word being in error, then the error correcting circuitry can cover the loss of the column without losing the entire memory array. Alternately, spare sense amplifiers may be included in each block of sense amplifiers together with the gates needed to insert a spare in place of a failed amplifier.

In the above-described embodiments of the present invention, the error correction circuit communicates the identity of the bits that are found to be erroneous to the controller as the various data words are delivered to the CPU. The time to communicate this information and have the controller store it can be longer than a memory cycle. Accordingly, some error events will be lost in the sense that the controller will not record and act on the events even though the error correction circuit corrected the errors. This situation is most likely when the CPU is executing a section of code that concentrates heavily on memory accesses in one region of the memory and that region has a weak storage cell. In principle, this situation can lead to the loss of information about a second weak storage cell in another region of memory that is not being accessed as often, because the controller misses the error report for the second weak storage cell because of an overload of error reports for the first weak storage cell.

One method for avoiding such lost reports is to utilize a buffer in error correction circuit that accumulates errors by the address at which the error occurred. Such a buffer can be constructed in a manner analogous to a CAM. A CAM stores a data record associated with a given address. Each time a request is made for data at a given address, the CAM recovers the

delays associated with moving data from disk to memory. In conventional memory systems, a block of bad memory in the middle of the RAM will cause the memory to fail if an address in the bad block is used. Hence, even if the system is smart enough to spot the bad RAM and reduce the available RAM size such that addresses above the bad address are used, the system will have lost half its memory. In effect, the system has to throw out the good RAM that is above the bad block of memory since the operating system assumes that memory is contiguous. In contrast, the present invention “moves” the bad addresses to the back of the memory leaving all of the good memory in what appears to be a continuous block. Hence, even if a memory according to the present invention runs out of spare rows, it will always be able to use the good rows.

As noted above, the internal organization of a memory according to the preferred embodiment of the present invention provides for much faster testing than is available in conventional memories. Conventional memories are tested by storing and reading data across the memory bus between the central processor and the memory. Hence, the data rates are limited by the width and speed of the system bus which is typically less than or equal to 64 bits wide and which runs at approximately 100 MHz. Accordingly, periodic exhaustive or even thorough testing of a memory for cells to detect intermittent failures is not possible in prior art memory designs when memory size begins to exceed 100 Mbytes.

As a result of these limitations, memories are tested thoroughly once after fabrication and repairs made to defective parts, if possible. This process requires specialized test equipment and long testing times. Both of these factors increase the cost of the final memory. In addition, even with specialized test equipment, the time available for testing does not permit tests to be run over a full range of possible operating conditions. For example, a part that does not operate at the highest possible clock frequency may operate satisfactorily at a lower clock frequency, and hence, be useable. The cost of identifying such parts, however, precludes their reclamation and sale. In any event, once memory passes, it is not thoroughly tested again until a system in which it is located fails and diagnostic procedures are instituted to determine the cause of failure.

As noted above, the present invention overcomes these difficulties by providing a fast highly parallel system of memory testing on the memory chip. This aspect of the present invention will now be explained in more detail with the aid of Figure 3, which is a block diagram of a memory 100 according to the present invention. Memory 100 differs from memory 10 discussed above in two significant aspects. First, memory 100 is constructed from a plurality of memory banks. Exemplary memory banks are shown at 101 and 102. Each memory bank is similar to that shown in Figure 2. For example, memory bank 101 includes a controller 105, row and columns CAMs 109, ECC generator 108, storage array 111, and ECC correction circuit 107. The memory banks connect to a bus 104, which includes the various signal lines used to connect the memory to the external memory bus. A memory block interface 103 is used to connect, or disconnect, memory block 101 from the system bus 104. Hence, controller 105 can run tests on storage array 111 at the same time controller 205 in memory block 102 is running tests on the storage array in that memory block. Accordingly, the test time is reduced by a factor equal to the number of memory blocks in the memory.

In addition to the improvements provided by dividing the memory into blocks for testing, the present invention also takes advantage of the higher speeds with which the internal buses in the memory operate. In general, the internal buses will support clock rates that are higher than the system bus clock rate because of the short distances between components within the memory. In the preferred embodiment of the present invention, interface 103 includes a phase locked frequency generator that generates the clock signals used during testing from a slower clock signal on bus 104. In the preferred embodiment of the present invention, a number of different clock rates are tested to determine the highest frequency at which each memory block will function.

Finally, as noted above, the storage arrays utilize much wider rows than those used in conventional storage arrays. In one embodiment of the present invention of the type shown in Figure 2, each row includes 72 Kbits and there are 4 memory blocks providing a total of 1 Gbyte of storage. Hence, this embodiment reads and writes storage words that are effectively 288 Kbits long. During testing all 4 blocks read and write in parallel. During normal operation, 64 bit

words are read and written from the appropriate block in the memory as determined by the address associated with the word. The row assigned to that address is read in parallel providing 1024 64 bit words together with 8 ECC bits for each 64-bit word. This data is stored in a register that is part of the sense amplifier array. The appropriate 72 columns are directed to the ECC

5 circuitry by a cross-connect switch as described above with reference to memory 10 discussed above. It should be noted that after reading out the requested word, the remaining words are still available in the output register. Hence, if a second read request, directed to another word in this row is received prior to the contents of this register being altered, that request can be serviced in a fraction of the access time required to provide the first word from the row.

10 A write operation is somewhat more complicated. To write data to one of the positions in the row, the entire row is first read out into the output register discussed above. The new data, together with its additional ECC bits, is then written into the appropriate location in that register, and the entire row is written back into the storage array. The same cross-connect switch used to reconfigure the memory can be used to direct the 64 bits being written to the appropriate location

15 in the register.

Since the present invention can perform an extensive memory test during the power up cycle of a computer, the expensive and time consuming memory testing normally performed on the production line can be omitted. In the preferred embodiment of the present invention, the

20 memory is tested during burn-in at a number of different clock speeds. The results of the testing can be utilized to sort the memory by performance characteristics. For example, memory chips that operate over a wider range of temperature and/or frequency can be sorted and sold for a premium.

25 The embodiments discussed above utilize a CAM to map the physical addresses within the memory chip, i.e., row and columns in the storage array, to the logical addresses utilized by the computer system in which the memory functions. However, it will be obvious to those skilled in the art from the preceding discussion that any form of mapping circuit can be utilized

30 for this function. In addition, the mapping circuit can include some form of non-volatile memory

such as EEPROM or FLASH in which the address mapping is stored when power is turned off. Such storage allows the memory to be configured once during burn in. This type of memory would mimic conventional memory, and hence, provide upward compatibility in systems that do not test memory prior to operation.

5

The above-described embodiments of the present invention store the error correcting bits on the same row as the data bits. However, it will be obvious to those skilled in the art from the preceding discussion that these bits can be stored in a separate storage array which is accessed via row select circuitry or directly via the word address. In the later case, the error correcting code storage space is not reconfigurable; however, savings in CAM hardware are achieved.

10

The above-described embodiments of the present invention store one bit per memory storage cell. That is, the data values read from the bit lines are either 1 or 0. However, the present invention may also be utilized in "multi-level" memories in which each storage cell stores a plurality of bits in the form of an analog value that is decoded by the sense amplifiers which now include some form of analog-to-digital converter. In this case, the single data values run from 0 to V , where $V > 1$.

15

20

Various modifications to the present invention will become apparent to those skilled in the art from the foregoing description and accompanying drawings. Accordingly, the present invention is to be limited solely by the scope of the following claims.

WHAT IS CLAIMED IS:

1. A reconfigurable memory comprising:

M bit lines, where $M > 1$;

a plurality of row lines;

an array of memory storage cells, each memory storage cell storing a data value and comprising circuitry for coupling that data value to one of said bit lines in response to a row control signal on one of said row lines;

a row select circuit for generating said row control signal on one of said row lines in response to a row address being coupled to said row select circuit, said row select circuit comprising a memory for storing a mapping of said row addresses to said row lines, said mapping determining which of said row lines is selected for each possible value of said row address; and

a controller for determining that one of said memory cells is defective and for altering said mapping to eliminate references to that row line that causes that defective storage cell to couple a data value to one of said bit lines.

2. The reconfigurable memory of Claim 1 wherein said controller tests all of said memory storage cells to determine if any of said memory storage cells is defective each time power is applied to said controller and wherein said controller eliminates references in said mapping to row lines that cause said detected defective storage cells to couple data values to said bit lines.

3. The reconfigurable memory of Claim 2 wherein said controller assigns a row address to each of said reference lines that was not eliminated because of a defective memory cell and

wherein said controller communicates the maximum number of rows available for storing data values after the elimination of said defective row references.

4. The reconfigurable memory of Claim 1 wherein said memory further comprises a single cell memory for storing a plurality of single data values, each data value corresponding to one of said row addresses and one of said bit lines and wherein said controller further comprises a circuit for causing that data value stored in said single cell memory for one of said row addresses and bit lines to replace that value stored in said memory storage cell coupled to that bit line when that row address is coupled to said row select circuit.

5. The reconfigurable memory of Claim 1 further comprising a word assembly circuit for selecting N bit lines from said M bit lines, where N is less than or equal to M, said word assembly circuit comprising a memory for storing a mapping specifying said N bit lines for each possible row address, wherein said controller alters said mapping to eliminate a reference in said mapping to a bit line that causes a defective storage cell to couple data to a bit line in response to one of said row addresses.

6. The reconfigurable memory of Claim 5 wherein said word assembly circuit comprises a cross-connect switch for coupling said M bit lines to N data lines.

7. The reconfigurable memory of Claim 6 wherein said cross-connect switch is partially populated such that only selected ones of said M bit lines can be connected to any particular data line.

8. The reconfigurable memory of Claim 1 further comprising an error correcting circuit for detecting errors in data words, said error correcting circuit generating a corrected data word and an error data word from N data values coupled thereto, said error data word indicating which of said N data values, if any, was erroneous; and

a word assembly circuit for connecting N of said M bit lines to said error correcting circuit, where N is less than or equal to M; wherein said control circuit is connected to said error correcting circuit and receives said error data words and said row addresses, said control circuit altering said mapping in said row select circuit in response to said error data words.

5

9. The memory of Claim 8 wherein $N < M$ and wherein said word assembly circuit comprises a cross-connect circuit for connecting N of said M bit lines to said error correcting circuit.

10

10. The memory of Claim 8 further comprising an error correcting code generating circuit for generating an error correcting data word from a data word, said error correcting data word being stored in a location corresponding to said data word in said memory.

15

11. The memory of Claim 10, wherein said error correcting data word is stored in the same row of memory cells as said data word used to generate said error correcting data word.

Dynamic Configuration of Storage Arrays

ABSTRACT

5

A reconfigurable memory having M bit lines and a plurality of row lines, where $M > 1$.

10

The memory includes an array of memory storage cells, each memory storage cell storing a data value. The data value is read from or into the storage cells by coupling that data value to one of the bit lines in response to a row control signal on one of the row lines. A row select circuit generates the row control signal on one of the row lines in response to a row address being coupled to the row select circuit. The row select circuit includes a memory for storing a mapping of the row addresses to the row lines that determines which of the row lines is selected for each possible value of the row address. The memory includes a plurality of sense amplifiers, one such sense amplifier being connected to each of the bit lines for measuring a signal value on that bit line. A controller that is part of the memory tests the memory storage cells both at power up and run time to detect defective memory storage cells. The controller uses an error correcting code scheme to detect errors during the actual operation of the memory. The memory includes sufficient spare rows and columns to allow the controller to substitute spares for rows or columns having defective memory storage cells.

15

20




FIGURE 1




FIGURE 2




FIGURE 3

100

DECLARATION FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

"Dynamic Configuration of Storage Arrays"

the specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, Section 1.56(a).

I hereby claim foreign priority benefits under Title 35, United States Code, § 119 of any foreign application(s) for patent or inventor's certificate listed below and so identified, and I have also identified below any foreign application for patent or inventor's certificate on this invention filed by me or my legal representatives or assigns and having a filing date before that of the application on which priority is claimed.

<u>Number</u>	<u>Country</u>	<u>Day/Month/Year Filed</u>	<u>Priority Claimed - Yes or No</u>
none			

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, § 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

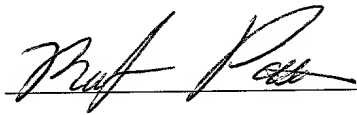
<u>Application Serial No.</u>	<u>Filing Date</u>	<u>Status</u>
none		

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint Calvin B. Ward (Reg. No. 30,896), with full power of substitution and revocation, to prosecute this application and to transact all business in the United States Patent and Trademark Office connected therewith and request that all correspondence and telephone calls in respect to this application be directed to The Law Offices of Calvin B. Ward, 18 Crow Canyon Court, Suite 305, San Ramon, CA 94583, Telephone No. (925) 855-0413.

Full name of inventor: Robert Patti

Inventor's signature:



Date:

5/23/00

Residence and Post Office Address:

1 South 751 Avon Drive
Warrenville, IL 60555

Citizenship:

U.S.A.